CLAIMS

1. In a cluster of computing nodes having shared access to one or more file systems in data storage using parallel file system software, a method for managing the data storage, comprising:

initiating a session of a data management application on a first one of the nodes;

receiving a request submitted to the parallel file system software at a second one of the nodes to mount one of the file systems in the data storage on the second one of the nodes; and

sending a mount event message from the second node to the first node responsive to the request, for processing by the data management application on the first node.

2. A method according to claim 1, and comprising mounting first and second instances of the one of the file systems on the first and second nodes, respectively, responsive to the mount event message.

3. A method according to claim 2, and comprising:

receiving a further request at the second node to unmount the second instance of the one of the file systems at the second node;

sending, responsive to the further request, a preunmount event message to the first node; and

responding to the preunmount event message so as to permit unmounting of the second file system instance without unmounting the first file system instance.

4. A method according to claim 3, wherein responding to the preunmount event message comprises determining at the

first node, responsive to one or more flags set in the preunmount event message, whether the request was submitted on the first node or on another one of the nodes.

5. A method according to claim 3, and comprising:

receiving the preunmount event message at the first node;

obtaining a data management access right from a physical file system (PFS) software component at the first node responsive to the preunmount event message; and

processing the preunmount event message using the access right.

6. A method according to claim 3, wherein receiving the request comprises receiving first and second requests to mount different ones of the file systems in the data storage, and wherein receiving the further request comprises receiving further first and second requests to unmount the different ones of the file systems, and wherein sending the preunmount event message comprises, responsive to dispositions set for the different ones of the file systems, sending a first preunmount event message to the first node responsive to the first unmount request, and sending a second preunmount event message responsive to the second unmount request to a further node, on which a further data management application session has been initiated.

7. A method according to claim 3, wherein responding to the preunmount event message comprises sending a reply to the message from the first node to the second node, and comprising, responsive to the reply, unmounting the

second file system instance and sending an unmount event message from the second node to the first node.

8. A method according to claim 7, and comprising determining at the first node, responsive to one or more flags set in the unmount event message, whether the further request was submitted on the first node or on another one of the nodes.

9. A method according to claim 1, and comprising determining at the first node, responsive to one or more flags set in the mount event message, whether the request was submitted on the first node or on another one of the nodes.

10. A method according to claim 1, wherein initiating the session comprises initiating the session in accordance with a data management application programming interface (DMAPI) of the parallel file system software, and wherein receiving the request and sending the mount event message comprise processing the request and sending the message using the DMAPI.

11. A method according to claim 10, and comprising receiving an unmount request to unmount the file system from the second node using the DMAPI, and sending a preunmount event message to the first node responsive to the unmount request using the DMAPI, for processing by the data management application on the first node.

12. A method according to claim 11, and comprising sending a reply to the preunmount event message from the first node to the second node using the DMAPI, and, responsive to the reply, unmounting the file system at

the second node, and sending an unmount event message to the first node using the DMAPI.

13. A method according to claim 10, and comprising receiving and processing the event message at the first node using one or more functions of the DMAPI called by the data management application.

14. A method according to claim 10, wherein sending the event message comprises sending the message for processing in accordance with a disposition specified by the data management application using the DMAPI for association with an event generated by the file operation.

15. A method according to claim 10, wherein sending the event message comprises setting one or more flags in the message to indicate whether the request was submitted on the first node or on another one of the nodes.

16. A method according to claim 10, and comprising invoking a function of the DMAPI to obtain mount information regarding the one of the file systems, and wherein in a response provided by the function, one or more flags are set to indicate whether the one of the file systems is mounted on the first node or on another one of the nodes in the cluster or on both the first node and on another one of the nodes in the cluster.

17. A method according to claim 1, and comprising:

receiving a response to the mount event message from the data management application on the first node; and

mounting an instance of the one of the file systems on the second node subject to the response from the data management application on the first node.

18. A method according to claim 1, and comprising receiving a further request submitted to the parallel file system software to mount the one of the file systems on a further one of the nodes, and sending a further mount event message responsive to the further request for processing by the data management application on the first node.

19. A method according to claim 18, wherein the further one of the nodes is the first node.

20. A method according to claim 19, and comprising receiving first and second unmount requests to unmount the file system from the second node and from the further one of the nodes, and generating first and second preunmount event messages at the second node and at the further one of the nodes responsive to the first and second unmount requests, for processing by the data management application on the first node.

21. A method according to claim 20, and comprising sending a reply to the first and second preunmount event messages from the first node to the second node and to the further one of the nodes, and, responsive to the reply, unmounting the file system at the second node and the further one of the nodes, and generating respective unmount event messages at the second node and at the further one of the nodes.

22. A method according to claim 1, wherein initiating the session of the data management application comprises initiating a data migration application, so as to free storage space on at least one of the volumes of data storage.

23. Computing apparatus, comprising:

one or more volumes of data storage, arranged to store data in one or more file systems; and

a plurality of computing nodes, linked to access the volumes of data storage using parallel file system software, and arranged so as to enable a data management application to initiate a data management session on a first one of the nodes, so that when a request is submitted to the parallel file system software at a second one of the nodes to mount one of the file systems in the data storage on the second one of the nodes, a mount event message is sent from the second node to the first node responsive to the request, for processing by the data management application on the first node.

24. Apparatus according to claim 23, wherein the nodes are arranged so that first and second instances of the one of the file systems are mounted on the first and second nodes, respectively, responsive to the mount event message.

25. Apparatus according to claim 24, wherein responsive to a further request at the second node to unmount the second instance of the one of the file systems at the second node, a preunmount event message is sent to the first node, which is arranged to respond to the preunmount event message so as to permit unmounting of the second file system instance without unmounting the first file system instance.

26. Apparatus according to claim 25, wherein the first node is arranged to respond to the unmount event message by determining, responsive to one or more flags set in

the preunmount event message, whether the request was submitted on the first node or on another one of the nodes.

27. Apparatus according to claim 25, wherein the first node is arranged, upon receiving the preunmount event message, to obtain a data management access right from a physical file system (PFS) software component at the first node responsive to the preunmount event message, and to process the preunmount event message using the access right.

28. Apparatus according to claim 25, wherein the request comprises first and second requests to mount different ones of the file systems in the data storage, and wherein the further request comprises further first and second requests to unmount the different ones of the file systems, and wherein the nodes are arranged, responsive to dispositions set for the different ones of the file systems, to send a first preunmount event message to the first node responsive to the first unmount request, and to send a second preunmount event message responsive to the second unmount request to a further node, on which a further data management application session has been initiated.

29. Apparatus according to claim 25, wherein the first node is arranged to send a reply to the message to the second node, and responsive to the reply, the second node is arranged to unmount the second file system instance and to send an unmount event message to the first node.

30. Apparatus according to claim 29, wherein the first node is arranged to determine, responsive to one or more

flags set in the unmount event message, whether the further request was submitted on the first node or on another one of the nodes.

31. Apparatus according to claim 23, wherein the first node is arranged to determine, responsive to one or more flags set in the mount event message, whether the request was submitted on the first node or on another one of the nodes.

32. Apparatus according to claim 23, wherein the session is initiated in accordance with a data management application programming interface (DMAPI) of the parallel file system software, and wherein the request is processed and the mount event message is sent using the DMAPI.

33. Apparatus according to claim 32, wherein when an unmount request is received to unmount the file system from the second node using the DMAPI, a preunmount event message is sent to the first node responsive to the unmount request using the DMAPI, for processing by the data management application on the first node.

34. Apparatus according to claim 33, wherein the first node is arranged to send a reply to the preunmount event message to the second node using the DMAPI, wherein responsive to the reply, the file system is unmounted at the second node, and an unmount event message is sent to the first node using the DMAPI.

35. Apparatus according to claim 32, wherein the event message is received and processed at the first node using one or more functions of the DMAPI called by the data management application.

36. Apparatus according to claim 32, wherein the mount event message is sent for processing in accordance with a disposition specified by the data management application using the DMAPI for association with the mount event.

37. Apparatus according to claim 32, wherein one or more flags are set in the event message to indicate whether the request was submitted on the first node or on another one of the nodes.

38. Apparatus according to claim 32, wherein the first node is arranged to invoke a function of the DMAPI to obtain mount information regarding the one of the file systems, and wherein in a response provided by the function, one or more flags are set to indicate whether the one of the file systems is mounted on the first node or on another one of the nodes in the cluster or on both the first node and on another one of the nodes in the cluster.

39. Apparatus according to claim 23, wherein after the mount event message is received at the first node, an instance of the one of the file systems is mounted on the second node subject to the response from the data management application on the first node.

40. Apparatus according to claim 23, wherein responsive to a further request submitted to the parallel file system software to mount the one of the file systems on a further one of the nodes, a further mount event message responsive to the further request is sent for processing by the data management application on the first node.

41. Apparatus according to claim 40, wherein the further one of the nodes is the first node.

42. Apparatus according to claim 41, wherein upon receiving first and second unmount requests to unmount the file system from the second node and from the further one of the nodes, first and second preunmount event messages are generated at the second node and at the further one of the nodes responsive to the first and second unmount requests, for processing by the data management application on the first node.

43. Apparatus according to claim 42, wherein the first node is arranged to send a reply to the first and second preunmount event messages to the second node and to the further one of the nodes, and wherein, responsive to the reply, the file system is unmounted at the second node and the further one of the nodes, and respective unmount event messages are generated at the second node and at the further one of the nodes.

44. Apparatus according to claim 23, wherein the data management application comprises a data migration application, for freeing storage space on at least one of the volumes of data storage.

45. A computer software product for use in a cluster of computing nodes having shared access to one or file systems in data storage, accessed using parallel file system software, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the computing nodes, cause a session of a data management application to be initiated on a first one of the nodes, and in response to a request submitted to the parallel file system software at a second one of the nodes to

mount one of the file systems in the data storage on the second node, cause the second node to send a mount event message to the first node, for processing by the data management application on the first node.

46. A product according to claim 45, wherein the instructions cause the nodes to mount first and second instances of the one of the file systems on the first and second nodes, respectively, responsive to the mount event message.

47. A product according to claim 46, wherein responsive to a further request at the second node to unmount the second instance of the one of the file systems at the second node, the instructions cause a preunmount event message to be sent to the first node, and cause the first node to respond to the preunmount event message so as to permit unmounting of the second file system instance without unmounting the first file system instance.

48. A product according to claim 47, wherein the instructions cause the first node to respond to the preunmount event message by determining, responsive to one or more flags set in the preunmount event message, whether the request was submitted on the first node or on another one of the nodes.

49. A product according to claim 47, wherein the instructions cause the first node, upon receiving the preunmount event message, to obtain a data management access right from a physical file system (PFS) software component at the first node responsive to the preunmount event message, and to process the preunmount event message using the access right.

50. A product according to claim 47, wherein the request comprises first and second requests to mount different ones of the file systems in the data storage, and wherein the further request comprises further first and second requests to unmount the different ones of the file systems, and wherein the instructions cause the nodes, responsive to dispositions set for the different ones of the file systems, to send a first preunmount event message to the first node responsive to the first unmount request, and to send a second preunmount event message responsive to the second unmount request to a further node, on which a further data management application session has been initiated.

51. A product according to claim 47, wherein the instructions cause the first node to send a reply to the message to the second node, and cause the second node, responsive to the reply, to unmount the second file system instance and to send an unmount event message to the first node.

52. A product according to claim 51, wherein the instructions cause the first node to determine, responsive to one or more flags set in the unmount event message, whether the further request was submitted on the first node or on another one of the nodes.

53. A product according to claim 45, wherein the instructions cause the first node to determine, responsive to one or more flags set in the mount event message, whether the request was submitted on the first node or on another one of the nodes.

54. A product according to claim 45, wherein the product comprises a data management application programming interface (DMAPI) of the parallel file system software, and wherein the request is processed and the mount event message is sent using the DMAPI.

55. A product according to claim 54, wherein when an unmount request is received to unmount the file system from the second node using the DMAPI, the instructions cause a preunmount event message to be sent to the first node responsive to the unmount request using the DMAPI, for processing by the data management application on the first node.

56. A product according to claim 55, wherein the instructions cause the first node to send a reply to the preunmount event message to the second node using the DMAPI, wherein responsive to the reply, the file system is unmounted at the second node, and an unmount event message is sent to the first node using the DMAPI.

57. A product according to claim 54, wherein the event message is received and processed at the first node using one or more functions of the DMAPI called by the data management application.

58. A product according to claim 54, wherein the event message is sent for processing in accordance with a disposition specified by the data management application using the DMAPI for association with an event generated by the file system.

59. A product according to claim 54, wherein one or more flags are set in the event message to indicate whether

the request was submitted on the first node or on another one of the nodes.

60. A product according to claim 54, wherein the instructions cause the first node to invoke a function of the DMAPI to obtain mount information regarding the one of the file systems, and wherein in a response provided by the function, one or more flags are set to indicate whether the one of the file systems is mounted on the first node or on another one of the nodes in the cluster or on both the first node and on another one of the nodes in the cluster.

61. A product according to claim 45, wherein after the mount event message is received at the first node, an instance of the one of the file systems is mounted on the second node subject to the response from the data management application on the first node.

62. A product according to claim 45, wherein responsive to a further request submitted to the parallel file system software to mount the one of the file systems on a further one of the nodes, a further mount event message responsive to the further request is sent for processing by the data management application on the first node.

63. A product according to claim 62, wherein the further one of the nodes is the first node.

64. A product according to claim 63, wherein upon receiving first and second unmount requests to unmount the file system from the second node and from the further one of the nodes, the instructions cause first and second preunmount event messages to be generated at the second node and at the further one of the nodes responsive to

the first and second unmount requests, for processing by the data management application on the first node.

65. A product according to claim 64, wherein the instructions cause the first node to send a reply to the first and second preunmount event messages to the second node and to the further one of the nodes, and wherein, responsive to the reply, the file system is unmounted at the second node and the further one of the nodes, and respective unmount event messages are generated at the second node and at the further one of the nodes.

66. A product according to claim 45, wherein the data management application comprises a data migration application, for freeing storage space on at least one of the volumes of data storage.